

Brief, But Not Prolonged, Pauses in the Firing of Midbrain Dopamine Neurons Are Sufficient to Produce a Conditioned Inhibitor

Chun Yun Chang,¹ Matthew P.H. Gardner,¹ Jessica C. Conroy,¹ Leslie R. Whitaker,¹ and  Geoffrey Schoenbaum^{1,2,3}

¹National Institute on Drug Abuse Intramural Research Program, Cellular Neurobiology Research Branch, Behavioral Neurophysiology Research Section, Baltimore, Maryland 21224, ²Department of Anatomy and Neurobiology, University of Maryland, School of Medicine, Baltimore, Maryland 21201, and

³Solomon H. Snyder Department of Neuroscience, The Johns Hopkins University, Baltimore, Maryland 21287

Prediction errors are critical for associative learning. In the brain, these errors are thought to be signaled, in part, by midbrain dopamine neurons. However, although there is substantial direct evidence that brief increases in the firing of these neurons can mimic positive prediction errors, there is less evidence that brief pauses mimic negative errors. Whereas pauses in the firing of midbrain dopamine neurons can substitute for missing negative prediction errors to drive extinction, it has been suggested that this effect might be attributable to changes in salience rather than the operation of this signal as a negative prediction error. Here we address this concern by showing that the same pattern of inhibition will create a cue able to meet the classic definition of a conditioned inhibitor by showing suppression of responding in a summation test and slower learning in a retardation test. Importantly, these classic criteria were designed to rule out explanations founded on attention or salience; thus the results cannot be explained in this manner. We also show that this pattern of behavior is not produced by a single, prolonged, ramped period of inhibition, suggesting that it is precisely timed, sudden change and not duration that conveys the teaching signal.

Key words: associative learning; conditioned inhibition; dopamine; prediction error; rat; ventral tegmental area

Significance Statement

Here we show that brief pauses in the firing of midbrain dopamine neurons are sufficient to produce a cue that meets the classic criteria defining a conditioned inhibitor, or a cue that predicts the omission of a reward. These criteria were developed to distinguish actual learning from salience or attentional effects; thus these results formally show that brief pauses in the firing of dopamine neurons can serve as key teaching signals in the brain. Interestingly, this was not true for gradual prolonged pauses, suggesting it is the dynamic change in firing that serves as the teaching signal.

Introduction

Prediction errors are thought to be responsible for associative learning (Rescorla and Wagner, 1972; Sutton, 1988), and single-unit (Mirenowicz and Schultz, 1994; Waelti et al., 2001; Pan et al.,

2005; Roesch et al., 2007), imaging (D'Ardenne et al., 2008), and voltametry (Day et al., 2007; Hart et al., 2014) studies have firmly established a correlative link between transient changes in dopamine neuron activity and reward prediction errors. Furthermore, the development of optogenetic means of directly manipulating the firing of these neurons in real time has yielded causal evidence supporting this linkage. Specifically, brief changes in the firing of these neurons can mimic and sometimes disrupt the operation of positive and negative prediction errors to drive learning about rewarding and even nonrewarding events (Frank et al., 2007; Tsai et al., 2009; Zweifel et al., 2009; Kim et al., 2012; Steinberg et al., 2013; Stopper et al., 2014; Chang et al., 2016, 2017; Sharpe et al., 2017).

However, alternative explanations for these data (particularly the causal evidence) abound, including most prominently the idea that learning driven by artificially induced transients is

Received Jan. 19, 2018; revised Aug. 10, 2018; accepted Aug. 20, 2018.

Author contributions: C.Y.C., L.R.W., and G.S. designed research; C.Y.C., M.P.H.G., J.C.C., and L.R.W. performed research; M.P.H.G. and J.C.C. contributed unpublished reagents/analytic tools; C.Y.C., L.R.W., and G.S. analyzed data; C.Y.C. and G.S. wrote the paper.

This work was supported by the Intramural Research Program at the National Institute on Drug Abuse (NIDA). We thank Dr. Karl Deisseroth and the Gene Therapy Center at the University of North Carolina at Chapel Hill for providing viral reagents and Dr. Brandon Harvey and the NIDA Optogenetic and Transgenic Core for their assistance. We also thank the NIDA *ex vivo* electrophysiology core for its assistance. The opinions expressed in this article are the authors' own and do not reflect the view of the NIH/Department of Health and Human Services.

The authors declare no competing financial interests.

Correspondence should be addressed to either of the following: Chun Yun Chang at the above address, E-mail: tina.chang@nih.gov; or Geoffrey Schoenbaum at the above address, E-mail: geoffrey.schoenbaum@nih.gov.

DOI:10.1523/JNEUROSCI.0144-18.2018

Copyright © 2018 the authors 0270-6474/18/388822-09\$15.00/0

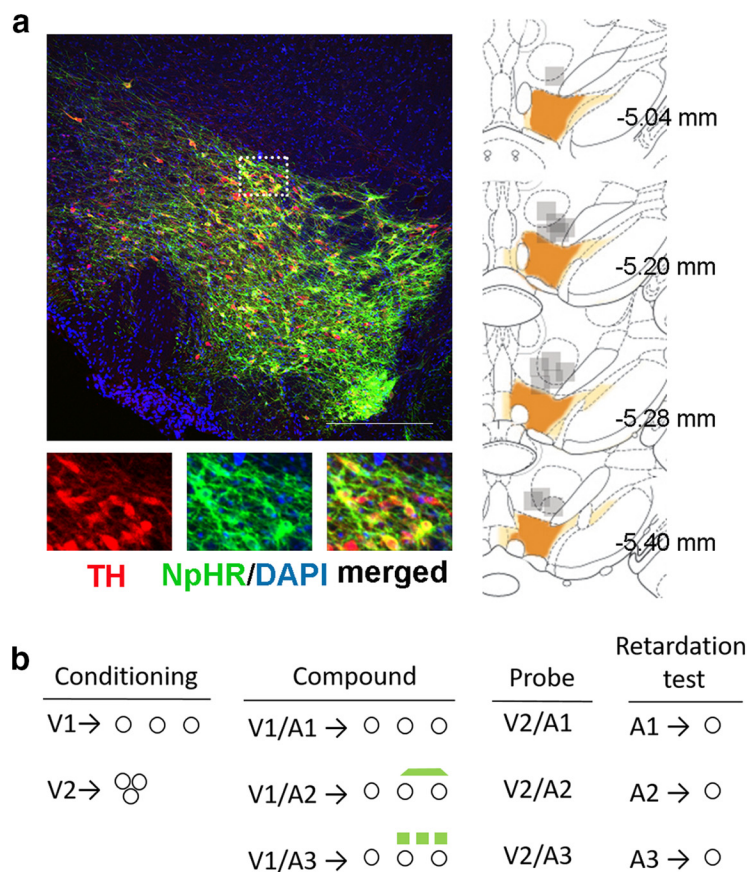


Figure 1. NpHR-eYFP expression, fiber implant positions, and task design in Experiment 1. **a**, NpHR-eYFP expression in TH-positive neurons in the VTA. Small photomicrographs illustrate the white boxed area in the main picture, showing TH staining (red, left), NpHR expression (middle, green), and coexpression (yellow, right). Scale bar, 300 μ m. The schematic to the far right illustrates the spread of NpHR-eYFP expression in the VTA and the localization of fiber tips; solid brown indicates the narrowest spread we observed across rats, light brown indicates the widest spread, and the gray squares show the fiber tip positions. **b**, Illustration of task design; see the text for a full description. V1 and V2 indicate solid and flashing panel lights, counterbalanced; A1, A2, and A3 indicate arduino-based melodies, counterbalanced; the circle indicates a 45 mg plain sucrose pellet; and green shows the approximate timing and duration of light delivery.

caused by nonspecific effects on attention rather than the mimicry of an endogenous error signal. For example, evidence that brief pauses in the firing of dopamine neurons can mimic or substitute for a missing negative prediction error to support extinction learning might instead be attributable to diminished salience or distraction. This idea is consistent with suggestions that the firing of dopamine neurons can signal novelty or salience (Kakade and Dayan, 2002; Bromberg-Martin et al., 2010; Schultz, 2016). Addressing these criticisms is important because it is the full pattern of the firing of these neurons that caused them to be so closely equated with prediction errors. Thus, in the causal studies, it is critical that dopamine transients operate like prediction errors in each direction (and in other ways).

With this in mind, here we address the possibility that brief pauses in the firing of midbrain dopamine neurons cause learning because of effects on salience rather than because of their ability to serve as negative prediction errors. To do this, we assess whether the same pattern of artificial inhibition that supports extinction learning can also turn a cue into a conditioned inhibitor. A conditioned inhibitor is a cue that predicts the nonoccurrence of events, typically a reward, predicted by other cues (Rescorla, 1969; Rescorla and Holland, 1977). It is normally created by pairing the cue that is to become a conditioned inhibitor with a reward-predicting cue and then omitting that reward. If

this is done repeatedly, the added cue becomes conditioned to predict the omission of that reward. Classically, this ability is demonstrated in a (negative) summation test, in which the conditioned inhibitor must suppress responding to another reward-predicting cue, combined with a retardation of learning test, in which the conditioned inhibitor must exhibit slower conditioning to predict the reward. Only cues that meet both criteria are considered to show conditioned inhibition, because only both effects can rule out explanations based on changes in salience. For example, a putative conditioned inhibitor might cause (negative) summation because it is highly salient or distracting if it is viewed as predictive of uncertain reward (Pearce and Hall, 1980). However, if this is the case, then the cue should facilitate learning in the retardation test since salient cues are more associable (Pearce and Hall, 1980). Conversely the putative conditioned inhibitor might retard learning if it has been made less salient by repeated presentation without reward (Mackintosh, 1975), but such an irrelevant or ignored cue should not reduce responding in the summation test. Thus it is the pattern of behavior in the pair of probe tests that defines conditioned inhibition.

Notably, this pattern of behavior will also rule out salience as an explanation of why brief pauses in the firing of dopamine neurons mimic negative prediction errors. If they act by altering salience, then we might obtain the expected result in one test, but we should fail to see the expected result in the other. In addition, we also

manipulated the firing of the dopamine neurons in two different ways to test whether it is the brief, appropriately timed, dynamic change in firing of these neurons that signals the error (Chang et al., 2016, 2017; Hamid et al., 2016) or whether the duration of the pause carries the information (Daw et al., 2002; Bayer and Glimcher, 2005; Bayer et al., 2007; Glimcher, 2011). We found that only sudden, brief pauses in the firing of midbrain dopamine neurons were sufficient to produce a cue that meets the classic criteria for conditioned inhibition; a single, equally long, ramped pause did not have this effect.

Materials and Methods

Subjects. Twenty-six transgenic rats (eight male and six female for the experiment in Figs. 1 and 2, three male and five female for the experiment in Figs. 4 and 5, and one male and three female for the experiment in Fig. 3) that carried a tyrosine hydroxylase (TH)-dependent Cre-expressing system on a Long-Evans background (NIDA animal breeding facility) served as subjects (Witten et al., 2011). Inclusion of gender as a factor in our analyses revealed no main effects or any interactions, so we collapsed this factor in reporting the data here (F values <0.66 ; p values >0.73). The rats were maintained on a 12 h light/dark cycle with unlimited access to food and water, except during the behavioral experiment when they were food restricted to maintain 85% of their baseline weight. All experimental procedures were conducted in accordance with the guidelines of

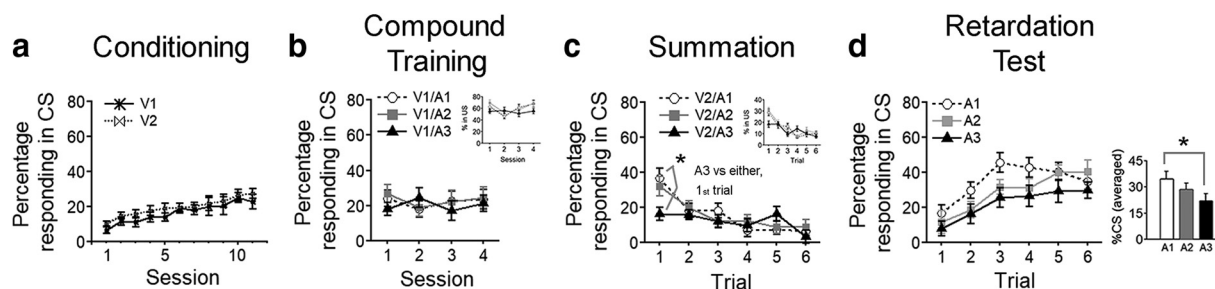


Figure 2. Brief, but not prolonged, pauses in the firing of midbrain dopamine neurons are sufficient to produce a conditioned inhibitor. Shown is the percentage of time the rats spent in the food cup on different trial types during the CS (main graphs) or US (insets) periods in conditioning (**a**), compound training (**b**), summation (**c**), and retardation probe testing (**d**). Green light, appropriate for activating NpHR to cause pauses in dopamine neuron firing, was delivered during delivery of the second and third food pellets on A2 and A3 trials in compound training, as described in the main text and Figure 1*b*. Error bars indicate SEM. *Critical planned comparisons that were significant at $p < 0.05$ or better; see the text.

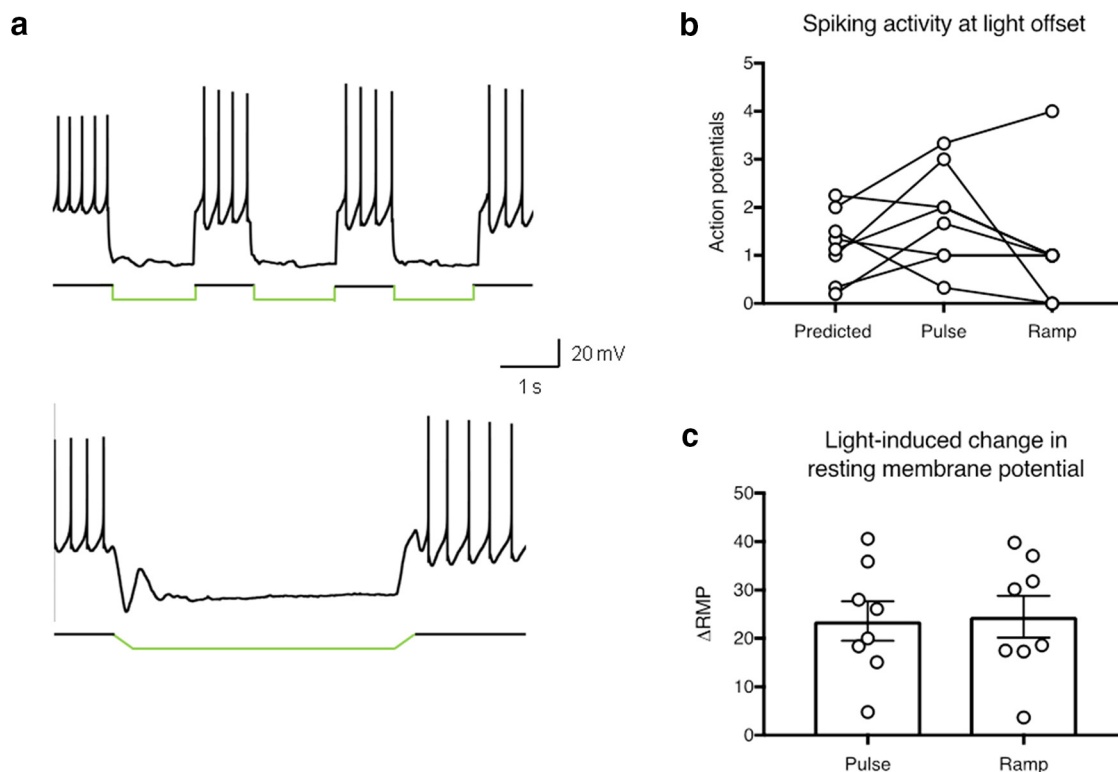


Figure 3. Effects of brief versus prolonged light delivery on spiking activity and membrane potential of dopamine neurons recorded *ex vivo*. **a**, Example traces show the effect of brief square pulses of light (top) and prolonged, ramped light delivery (bottom) on spiking activity in VTA dopamine neurons. **b**, Summary graph showing the number of spikes in the 500 ms after light offset compared with the number of spikes predicted from baseline activity in each recorded neuron. **c**, Light-induced hyperpolarization of the resting membrane potential after brief pulses or prolonged, ramped light delivery.

the Institutional Animal Care and Use Committee of the National Institutes of Health.

Surgical procedures. Rats (>275 g) received bilateral infusions of AAV5-EF1 α -DIO-NpHR3.0-eYFP into the ventral tegmental area [VTA; AP, -5 mm (referenced to bregma); ML, ± 0.7 mm (referenced to the midline); DV, -7.2 and -8 mm for male and -6.5 and -7.3 mm for female (referenced to the brain surface)]. Virus was obtained from the University of North Carolina at Chapel Hill Gene Therapy Center, courtesy of Dr. Karl Deisseroth. A total of 1 – 1.5 μ l of virus with a titer of $\geq 10^{12}$ vg/ml was injected at the rate of 0.1 μ l/min per injection site. The rats also were implanted with optic fibers bilaterally (200 μ m diameter, Thorlab; AP, -5.3 mm; ML, ± 2.61 mm; DV, -6.8 mm for male and -6.2 mm for female at 15° angle pointing to the midline).

Apparatus. Training was conducted in 10 standard behavioral chambers from Coulbourn Instruments, each enclosed in a sound-resistant shell. A food cup was recessed in the center of one end wall. Entries were monitored by a photobeam. A food dispenser containing 45 mg sucrose

pellets (plain, Bio-SERV) allowed delivery of pellets into the food cup (Graphic State 3, Coulbourn Instruments). Rats were shaped to retrieve pellets from the food cup before any conditioning.

Conditioned inhibition. Training consisted of three phases: conditioning, compound training, and probe testing. Conditioning consisted of 11 sessions. In each session, the rats were presented with two distinct 10 s visual cues (V1 and V2; 6 W light bulb with on/off patterns consisting of 0.5 s on and 0.5 s off for a total of 10 s, or 10 s on, counterbalanced) paired with three plain sucrose pellets, delivery of which commenced at the end of each cue. On V1 trials, pellets were delivered with 4 s spacing. On V2 trials, three pellets were delivered without any spacing. Each session consisted of 12 trials of V1 and 12 trials of V2, arranged in 6-trial blocks, the order of which varied from day to day for each rat. Trials were separated by intertrial intervals that varied randomly between 120, 130, and 140 s. After conditioning, the rats began compound training. Before the first session, the rats received six presentations each of three novel 10 s auditory cues (A1, A2, and A3; ~ 75 dB, customized Arduino-based melodies,

counterbalanced). Compound training consisted of four sessions. Trial structure was generally the same as during conditioning; each session consisted of three reminder trials of V1 alone, followed by six trials of V1 in compound with each auditory cue (V1/A1, V1/A2, and V1/A3). In the first experiment, each compound cue was followed by the same amount of reward used during conditioning; in the second experiment, the second and third pellets were omitted on V1/A1 trials. During the compound trials, light (~ 16 – 18 mW, Shanghai Laser & Optics Century) was delivered to the VTA, bilaterally, on V1/A2 and V1/A3 trials, timed to start 0.5 s before delivery of the second pellet. In Experiment 1, a green light (532 nm) was delivered; in Experiment 2, a blue light (473 nm) was delivered. On V1/A2 trials, a single 5.5 s pulse of light was delivered, which began and ended with a 0.5 s ramp to avoid sudden changes in firing of the dopamine neurons at either onset or offset of the light. On V1/A3 trials, three 1.5 s pulses of light were delivered, spaced by 1 s, without ramping, to mimic the sudden changes observed in dopamine neuron firing in response to naturally induced negative prediction errors. Note that the latter pattern is identical to that used in our prior study to test whether dopamine pauses could function as negative errors (Chang et al., 2016). There was no light delivered on V1/A1 trials. After the conclusion of compound training, the rats underwent a summation and then a retardation probe test. Trial structure was again generally the same as during conditioning except as noted below. In the summation test, the auditory cues were each presented six times, in compound with V2, in random order and without any reward. In the retardation test, the auditory cues were each presented six times, alone, followed by delivery of one sucrose pellet. For this test, cue presentation was blocked, with the cue order counterbalanced across rats as much as was practical, and learning about each cue was separated by a 2.5 h timeout period.

Response measures and statistics. The primary measure of conditioning was the percentage of time that each rat spent with its head in the food cup during each conditioned stimulus (CS) presentation before food delivery, as indicated by disruption of the photocell beam. Because the cues were presented in a block design in some parts of the experiment, we subtracted responding during a 10 s pre-CS baseline. Food cup responding is strongest at the end of the cue, so for clarity we also excluded the first half of the cue period from the data presented in the figures. However a direct comparison of the data shown with data from the initial 5 s period showed no interaction of period with any of the meaningful effects described in the main text (F values < 6.16 ; p values > 0.062). In addition, in some periods we also measured the amount of time the rats spent in the food cup during the post-CS period starting at the time of the first food pellet delivery. This was done to test for any aversive or distracting effect of the optogenetic manipulation during blocking or as an additional assessment of conditioned responding after the cue (in the absence of food) in the final probe test. Statistical comparisons were made by multifactor ANOVA (repeated-measures ANOVA), followed by planned comparisons as warranted by our hypothesis; unexpected effects were explored using a Bonferroni's correction (STATISICA, Statsoft, TIBCO Software).

Histology and immunohistochemistry. Rats that received viral infusions and fiber implants were killed with carbon dioxide and perfused with $1 \times$ PBS, followed by 4% paraformaldehyde (Santa Cruz Biotechnology). Fixed brains were cut in $40 \mu\text{m}$ sections to examine fiber tip position under a fluorescence microscope (Olympus Microscopy). For immunohistochemistry, the brain slices were first blocked in 10% donkey serum made in 0.1% Triton X-100/ $1 \times$ PBS and then incubated in anti-TH antisera (1:600, EMD Millipore), followed by Alexa Fluor 568 secondary antisera (1:1000, Invitrogen). The image of brain slices was acquired by a confocal microscope (Olympus FluoView 1000, America) and later analyzed in Adobe Photoshop. The VTA, including anterior (rostral and parabrachial pigmental area) and posterior (caudal, parabrachial pigmental area, paranigral nucleus, and medial substantia nigra pars medialis), of brain slices from AP -5.1 to -5.9 mm were analyzed from five subjects in Experiment 1 and from three subjects in Experiment 2. This encompasses the location targeted by our fibers and is likely to achieve good light penetration. For quantification, the intensity of three to four random $40 \times 40 \mu\text{m}$ square areas from the background were averaged to provide a baseline, and positive staining was defined as a signal 2.5 times

this baseline intensity, with a cell diameter larger than $5 \mu\text{m}$, colocalized within cells reactive to DAPI staining.

Ex vivo brain slice electrophysiology. TH-Cre rats that had received AAV-DIO-NpHR3.0-eYFP (NpHR) infused into VTA were deeply anesthetized with isoflurane and transcardially perfused with an ice-cold solution containing (in mM) 93 NMDG (N-Methyl-D-glucamine), 93 HCl, 2.5 KCl, 1.2 NaH_2PO_4 , 30 NaHCO_3 , 20 HEPES, 25 glucose, 5 Na-ascorbate, 2 thiourea, 3 Na-pyruvate, 10 MgSO_4 , and 0.5 CaCl_2 . Brains were rapidly removed, and horizontal midbrain slices ($220 \mu\text{m}$) were made using a Vibratome (Leica VT-1000S). Slices were then incubated in a holding solution containing (in mM) 92 NaCl, 2.5 KCl, 1.2 NaH_2PO_4 , 30 NaHCO_3 , 20 HEPES, 25 glucose, 5 Na-ascorbate, 2 thiourea, 3 Na-pyruvate, 2 MgSO_4 , and 0.5 CaCl_2 (32 – 34°C) for 15–30 min. After this, the holding chamber was kept at room temperature for the duration of the experiment. Slices were transferred to a recording chamber and superfused (2 – 3 ml/min) with artificial CSF containing (in mM) 126 NaCl, 2.5 KCl, 1.2 MgCl_2 , 2.4 CaCl_2 , 1.2 NaH_2PO_4 , 21.4 NaHCO_3 , and 11.1 glucose maintained at 32 – 34°C . All solutions were continually oxygenated (95% oxygen, 5% carbon dioxide). Glass pipettes (tip resistance, 2 – 4 M Ω) were used for whole-cell recordings and were filled with an intracellular solution containing (in mM) 115 K-gluconate, 20 KCl, 1.5 MgCl_2 , .025 EGTA, 10 HEPES, 2 Mg-ATP, 0.2 Na-GTP, and 10 Na_2 -phosphocreatine (pH 7.2–7.3, ~ 290 mOsm/kg).

Whole-cell current-clamp recordings were performed in visually identified neurons in the VTA. Virus-infected (eYFP+) cells were identified using scanning disk confocal microscopy (Olympus FV1000), and differential interference contrast optics were used to patch neurons. Light was delivered to stimulate NpHR *in vitro*, at an intensity of 16 – 18 mW, using the same equipment and two protocols (sudden and brief versus gradual and prolonged) used in the two behavioral experiments. The baseline firing rate was recorded in zero-current mode in the absence of light delivery.

Electrophysiology data were acquired using an Axopatch 200 B amplifier (Molecular Devices) in either voltage-clamp or current-clamp mode. Axograph X software (Axograph Scientific) was used to record and collect the data, which were filtered at 10 kHz and digitized at 4 – 20 kHz. Series resistance (R_s) was monitored with an injection of a hyperpolarizing current (-20 pA, 500 ms), and data were excluded if R_s changed $> 20\%$ during data acquisition. Recordings were discarded if series resistance or input resistance changed $> 10\%$ throughout the course of the recording.

Results

Brief, but not prolonged, pauses in the firing of midbrain dopamine neurons are sufficient to produce a conditioned inhibitor

Optogenetic and transgenic approaches used to target dopamine neurons in the VTA were similar to those used in prior studies (Chang et al., 2016, 2017; Sharpe et al., 2017). Briefly, 14 rats expressing Cre-recombinase from the TH promoter served as subjects (Witten et al., 2011). Each rat received bilateral infusions of NpHR within VTA, and fiber optics were implanted targeting this area bilaterally (Fig. 1a). Postmortem immunohistochemical verification showed a high degree of colocalization between Cre-dependent NpHR expression and TH in VTA; $\sim 90\%$ of NpHR-expressing cells in the VTA (307 from 344 cells counted in five rats) were immunoreactive to anti-TH antisera. After a 2 week period to allow viral expression, the rats were food deprived to 85% of their baseline body weight, and then we began training on the conditioned inhibition task.

Conditioned inhibition consisted of conditioning and compound training, followed by probe testing for the classic signs of conditioned inhibition: summation of responding and retardation of learning (Fig. 1b; Rescorla, 1969; Rescorla and Holland, 1977). During conditioning, the rats were trained to associate two novel 10 s visual cues, V1 and V2, with delivery of three plain sucrose pellets. Pellets were spaced differently for the two cues to

ensure that the cues were discriminated, so that we could use them independently in later testing. As illustrated in Figure 2*a*, rats developed conditioned responding to both cues across sessions ($F_{(1,13)} = 11.85$; $p = 0.004$) and showed slightly but significantly more responding to V2 than to V1, consistent with the modest difference in the spacing of the food pellets ($F_{(10,130)} = 7.18$; $p = 0.001$).

After conditioning, the rats underwent compound training. In these sessions, one of the visual cues, V1, was presented in compound with each of three 10 s auditory cues, A1, A2, and A3. A1 served as a control cue. It was presented in compound with V1, after which three sucrose pellets were delivered as expected, without any optogenetic manipulation. A2 and A3 were experimental cues. They were each presented in compound with V1, followed by three sucrose pellets, and we delivered green light into VTA (532 nm) to inhibit the dopamine neurons during pellet delivery. On A3 trials, we delivered three pulses of light, each lasting 1.5 s with a 1 s interval between each. This pattern was intended to mimic the sudden brief pauses in firing used in our prior study to substitute for negative prediction errors in extinction learning (Chang et al., 2016). This may not precisely mimic what actually would happen on omission of these rewards, though we have observed somewhat discrete pauses in dopamine neuron firing when expected rewards with a spacing of 0.5 s are omitted (Roesch et al., 2007). On A2 trials, we delivered a single pulse of light, lasting 5.5 s, including 0.5 s ramps at onset and offset. This pattern was intended to duplicate the total period of inhibition used on A3 trials, with gradual onset to avoid inducing the sudden changes in firing observed in response to naturally induced negative prediction errors (Mirenowicz and Schultz, 1994; Waelti et al., 2001; Pan et al., 2005; Roesch et al., 2007) and gradual offset to prevent rebound activity that has been observed after prolonged activation of opsins such as NpHR (Raimondo et al., 2012; Chuong et al., 2014; Mahn et al., 2016; Wiegert et al., 2017). Importantly, testing of these two protocols in an *ex vivo* brain slice preparation in a separate group of TH-Cre rats that had received NpHR infused into VTA confirmed that both caused similar effects on dopamine neuron firing and membrane potential ($t = 1.09$, $p = 0.31$) while the light was on, and there was no evidence of a significant rebound at the end of either the brief or prolonged duration pauses ($F_{(2,14)} = 1.65$; $p = 0.23$; Fig. 3).

As illustrated in Figure 2*b*, the addition of the auditory cues and the optogenetic manipulations had little effect on established responding. The rats continued to respond to V1 at about the same level observed in conditioning and this did not change across sessions ($F_{(3,39)} = 0.45$; $p = 0.72$), and there were no differences in responding on trials involving the different auditory cues, either during cue presentation ($F_{(2,26)} = 0.69$; $p = 0.51$) or when the food pellets themselves were delivered (Fig. 2*b*, inset; $F_{(2,26)} = 0.92$; $p = 0.41$). The lack of any differences in food cup responding, even during food delivery, suggests that briefly inhibiting the dopamine neurons was neither aversive nor generally distracting.

After compound training, the rats underwent the summation and retardation probe tests. In the summation test, each auditory cue was presented six times in compound with V2, the other visual cue that was conditioned earlier. No reward was delivered. As illustrated in Figure 2*c*, responding to V2 began relatively high and then declined quickly across the nonrewarded trials. Against this backdrop, however, responding to V2 was lower when it was paired with A3 than when it was paired with either A1 or A2. This effect was immediately and primarily evident on the first trial, and it was present both during the cue itself and in the period

after the cue when food was normally delivered. These notions were supported by ANOVA (cue \times trial), which showed significant effects of trial [CS: $F_{(5,65)} = 9.62$, $p = 0.000001$; unconditioned stimulus (US): $F_{(5,65)} = 8.18$, $p = 0.000005$] and interactions between cue and trial (CS: $F_{(10,130)} = 2.60$, $p = 0.007$; US: $F_{(10,130)} = 2.13$, $p = 0.03$). Planned contrasts revealed the interactions between cue and trial stemmed from lower responding to V2/A3 on the first trial, compared with either of the other cues, in both cue ($p = 0.0004$ vs A1; $p = 0.02$ vs A2) and reward ($p = 0.02$ vs A1; $p = 0.02$ vs A2) periods. The single, prolonged, ramped period of inhibition did not have any effect on responding to V2/A2 ($p = 0.25$ vs A1 in the cue period and 0.78 in the reward period).

The results of the summation test are consistent with prior evidence that sudden, brief pauses in the firing of the dopamine neurons support learning (Steinberg et al., 2013; Chang et al., 2016), since A3 was able to cause (negative) summation. Furthermore, since there was no effect on A2 responding, the data suggest that whatever is happening is not attributable to the duration of the pause but requires sudden, dynamic shifts in the dopamine neuron firing. However, the results do not distinguish an error signaling mechanism from attentional changes. For this, we must conduct the retardation of learning test.

To test for retardation of learning, each auditory cue was presented six times alone, followed by delivery of a single sucrose pellet. As illustrated in Figure 2*d*, rats conditioned to all three auditory cues, but they conditioned most quickly to A1 and more slowly to the other cues, especially A3. ANOVA (cue \times trial) confirmed this impression, revealing significant main effects of trial ($F_{(5,65)} = 11.15$; $p = 0.001$) and cue ($F_{(2,26)} = 3.79$; $p = 0.04$), and a planned contrast analyses indicated there was lower responding to A3 than to A1 ($p = 0.03$). This result distinguishes between the two explanations for the effects of brief inhibition of the dopamine neurons on summation. Whereas elevated salience can explain (negative) summation, elevated salience would cause facilitated rather than retarded learning for A3, since salient cues are more associable. On the other hand, a negative prediction error would explain both effects, since a negative prediction error would cause A3 to become predictive of omission of the reward, leading to both spontaneously less responding in the presence of V2 and slower conditioning when paired with reward itself. Interestingly, A2 also seemed to condition a bit more slowly than A1, although responding to A2 was not statistically different from either A1 ($p = 0.59$, Bonferroni's test) or A3 ($p = 0.52$, Bonferroni's test). A2 showed no evidence of functioning as a conditioned inhibitor in the summation test, so even if learning for A2 were retarded, the pattern across both tests would still be inconsistent with operation of a prediction error.

Mere delivery of light at the time of reward, whether brief or prolonged, is insufficient to produce a conditioned inhibitor

Although changes in A2 are intriguing, they are also problematic, since we intended this cue to control for nonspecific effects of light delivery within subjects in our design. Whereas A2 did not impact performance in the critical summation test, the intermediate effect in the retardation test raises concerns with this approach. To address these concerns, we repeated the experiment in a second group of rats, delivering blue light inappropriate for activating NpHR on the V1/A2 and V1/A3 trials and omitting the second and third food pellets on V1/A1 trials to create a conditioned inhibitor (Fig. 4*b*).

Eight rats expressing Cre-recombinase from the TH promoter received bilateral infusions of NpHR within VTA and fiber optics

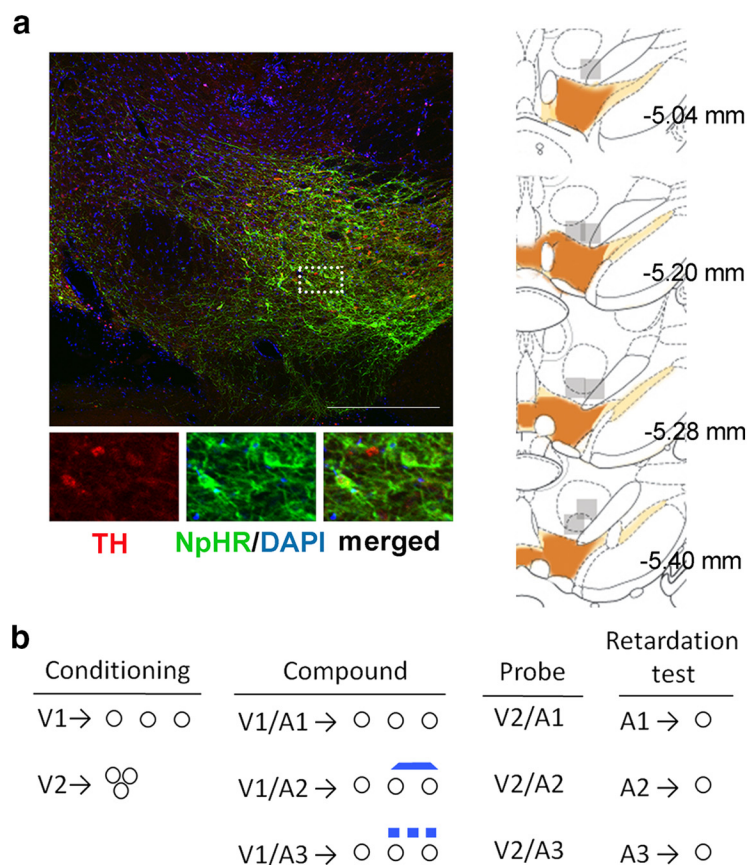


Figure 4. NpHR-eYFP expression, fiber implant positions, and task design in Experiment 2. **a**, NpHR-eYFP expression in TH-positive neurons in the VTA. Small photomicrographs illustrate the white boxed area in main picture, showing TH staining (red, left), NpHR expression (middle, green), and coexpression (yellow, right). Scale bar, 300 μ m. The schematic to the far right illustrates the spread of NpHR-eYFP expression in the VTA and the localization of fiber tips; solid brown indicates the narrowest spread we observed across rats, light brown indicates the widest spread, and the gray squares show the fiber tip positions. **b**, Illustration of task design; see the text for a full description. V1 and V2 indicate solid and flashing panel lights, counterbalanced; A1, A2, and A3 indicated arduino-based melodies, counterbalanced; the circle indicates the 45 mg plain sucrose pellet; and blue shows the approximate timing and duration of light delivery.

targeting this area bilaterally (Fig. 4a). One rat was excluded from the study after losing its implant during training; postmortem immunohistochemistry in the remaining subjects showed ~91% colocalization (92 of 101 cells counted from three rats) and viral spread and fiber tip localization similar to Experiment 1 (Fig. 4a).

After a 2 week period to allow viral expression, the rats were food deprived to 85% of their baseline body weight and underwent conditioning and compound training. As illustrated in Figure 5, *a* and *b*, the rats developed conditioned responding to both cues across sessions during conditioning ($F_{(10,60)} = 3.16$; $p = 0.002$) and learned to distinguish V1/A1 from the other two trial types, reflecting the reduced reward, an effect that was evident both during cue presentation ($F_{(2,12)} = 20.93$; $p = 0.0001$) and food pellet delivery (Fig. 5b, inset; $F_{(2,12)} = 79.51$; $p = 0.0000001$).

After compound training, the rats underwent the summation and retardation probe tests. In the summation test, illustrated in Figure 5c, responding to V2 was initially high in the presence of either A2 or A3 and declined as the test progressed without reward, whereas responding in the presence of A1 was immediately low, essentially at baseline. This pattern was present both during cue presentation and in the period after the cue when food was normally delivered. ANOVAs (cue \times trial) revealed significant

effects of trial (CS: $F_{(5,30)} = 2.99$, $p = 0.026$; US: $F_{(5,30)} = 4.58$, $p = 0.003$) and cue (CS: $F_{(2,12)} = 6.23$, $p = 0.01$; US: $F_{(2,12)} = 12.16$, $p = 0.001$), and planned contrasts indicated that responding to V2/A1 on the first trial was lower than to either of the other cues (CS: $p = 0.01$ vs A2 and 0.02 vs A3; US: $p = 0.04$ vs A2 and 0.01 vs A3), whereas responding to V2 in the presence of A2 and A3 did not differ (CS, $p = 0.15$; US, $p = 0.28$). In the retardation test, illustrated in Figure 5d, rats conditioned to all three auditory cues, but conditioning was less to A1 than to either of the other two cues. ANOVA (cue \times trial) revealed significant main effects of trial ($F_{(5,30)} = 5.54$; $p = 0.001$) and cue ($F_{(2,12)} = 4.92$; $p = 0.03$), with planned contrasts confirming that responding to A2 and A3 was similar ($p = 0.73$) and significantly higher than responding to A1 ($p = 0.01$ vs A1 and 0.03 vs A2).

Combined with the prior experiment, these results show delivery of light into VTA at the time of reward is insufficient to mimic effects of reward omission to create a conditioned inhibitor unless the light is of the proper wavelength to activate NpHR to suppress activity in the dopamine neurons.

Discussion

Here we have shown that sudden, brief pauses in the firing of midbrain dopamine neurons will produce a cue that meets the classic criteria for conditioned inhibition: (negative) summation and retardation of learning. Specifically, when a reward-predicting cue was presented with another cue, followed by reward plus several brief, transient pauses in the firing of VTA dopamine neurons, the added cue spontaneously suppressed responding to another cue predictive of the reward and was slower to condition to predict the reward. As discussed earlier, this pattern of results cannot be explained by changes in salience or associability, whereas it is consistent with prior claims that brief pauses in the firing of these neurons serve as negative prediction errors (Steinberg et al., 2013; Chang et al., 2016). Notably, a single, prolonged, gradual pause spanning the same period did not have this effect, supporting our prior use of a longer pause to assess whether dopamine transients are necessary for learning (Chang et al., 2017; Sharpe et al., 2017).

Although it was not intended to be the focus of the study, it is interesting to speculate on why a single, prolonged, gradual pause did not have these effects. Since we did not independently manipulate the number of pauses or their duration and speed of onset and offset, it is impossible to say which is the critical parameter. But the effects of the two protocols on the membrane potential and firing activity of dopamine neurons *in vitro* were remarkably similar, suggesting it was not because one protocol failed to inhibit the neurons or resulted in dramatically different patterns of firing. One possible explanation for the difference, which would only be evident across large populations, would be if the sudden onset of full-intensity light were more effective than

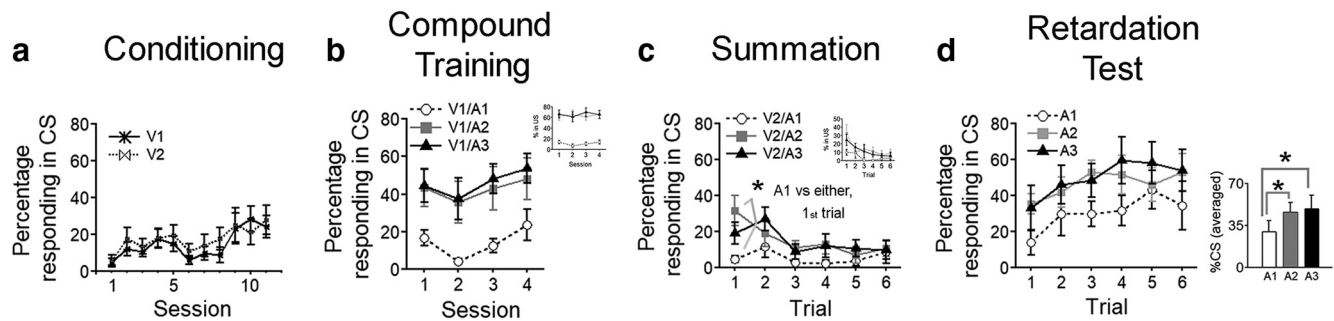


Figure 5. Mere delivery of light at the time of reward, whether brief or prolonged, is insufficient to produce a conditioned inhibitor. Shown is the percentage of time the rats spent in the food cup on different trial types during the CS (main graphs) or US (insets) periods in conditioning (**a**), compound training (**b**), summation (**c**), and retardation probe testing (**d**). Blue light, inappropriate for activating NpHR to cause pauses in dopamine neuron firing, was delivered during delivery of the second and third food pellets on A2 and A3 trials in compound training, as described in the main text and Figure 1*b*. Error bars indicate SEM. *Critical planned comparisons that were significant at $p < 0.05$ or better; see the text.

the ramped onset at promoting synchronized pauses across many neurons. This would theoretically occur if dopamine neurons, either because of natural (different proportions of inhibitory and excitatory input) or artificial (viral expression, location relative to fiber) factors, were prevented from spiking at different thresholds. If this were true, then the faster the light reaches its full intensity, the better synchronized in time the change in firing would be in the population. It has been shown that the effectiveness of dopamine neuron firing pauses in decreasing postsynaptic receptor occupancy is heavily dependent on the degree of synchrony within the dopamine neuron population; synchronous firing pauses across multiple neurons are much more effective in decreasing receptor occupancy than asynchronous activity decreases (Dreyer et al., 2010). Combined with suggestions that it is the dynamic change in dopamine firing that signals the error and not its absolute level (Hamid et al., 2016), this explanation provides a potential framework within which to understand the different effects of sudden brief versus gradual prolonged pauses.

Of course, we did observe a modest effect of gradual and more prolonged inhibition of the dopamine neurons on the ability of the cue to support new learning in the retardation test. Why would this occur if the gradual pause is ineffective as a teaching signal? One possibility is that it is an artifact of rebound excitation known to occur at the end of the long periods of inhibition (Raimondo et al., 2012). We believe several factors suggest this is not the case. First, light was delivered at full intensity for only 4.5 s, a period associated with relatively modest shifts in the reversal potential of the GABA_A receptor, which is thought to be responsible for rebound excitation after NpHR activation (Raimondo et al., 2012). Second, light offset was gradual, a manipulation thought to prevent or at least substantially mitigate rebound effects (Chuong et al., 2014; Mahn et al., 2016). Third, and perhaps most importantly, when the same protocol was used *in vitro* to inhibit dopamine neurons, there was no evidence of rebound. And finally, rebound excitation in the presence of reward should create a conditioned excitor (Steinberg et al., 2013). If A2 were a conditioned excitor, it would have caused positive summation and facilitation of learning, which it did not. Thus, we believe we can dispense with this explanation.

Another somewhat trivial possibility is that pairing A2 with light caused this cue to become distracting in the later retardation test. This might have affected A2 more than A3, either because of the longer period of constant light delivery or because of the absence of a meaningful competing physiological effect on the dopamine neurons. However, there was no evidence of distraction as a result of light delivery in compound training, and when

we repeated the experiment using light of the wrong wavelength, we found that A2 and A3 were treated identically and unlike conditioned inhibitors. This suggests that prolonged light alone cannot account for the effect.

This leaves the possibility that the retarded learning was the result of the prolonged reduction in dopamine neuron activity. Although this pattern of results cannot be explained by a prediction error mechanism, it would classically be interpreted as an effect on salience or even general value, independent of associative learning. Although clearly speculative, if this effect were found to be reliable, it would provide support for a dual-component account of dopamine function (Schultz, 2016), though at a time scale more consistent with proposals regarding tonic dopamine (Hamid et al., 2016).

On the other hand, our main result that brief pauses in the firing of dopamine neurons are sufficient to create a conditioned inhibitor provides little support for proposals that brief, transient changes independently convey information about salience or novelty (Kakade and Dayan, 2002; Bromberg-Martin et al., 2010; Schultz, 2016). This idea arose from correlative findings that are otherwise hard to fit into the popular account that dopamine transients signal cached-value prediction errors (Horvitz et al., 1997; Horvitz, 2000; Tobler et al., 2003; Matsumoto and Hikosaka, 2009). We have recently published data suggesting that dopamine transients may signal errors in the prediction of events more broadly (Takahashi et al., 2017). In this framework, apparent responses to novelty or salience can be understood as more general or informational error signals (Bromberg-Martin and Hikosaka, 2009; Gershman and Schoenbaum, 2017; Langdon et al., 2018). The failure to see any effect of transients here that would be consistent with salience is in keeping with this proposal.

An interesting question for the future is what information is learned under the influence of these putative negative prediction errors. Cached-value errors should not endow cues with the ability to predict specific events (Langdon et al., 2018). Thus, if the dopamine transient acts only as a cached-value error, the current results make sense, but only in a sort of arithmetic manner. In other words, the rats respond less when a cue is paired with inhibition of the dopamine neurons because it is less valuable, so that cue appears to act as a conditioned inhibitor if presented with a reward-predicting cue (e.g., $2 - 1 = 1$). But this should only work if the cue is tested with the same action that was present during learning, and it should not matter if the particular food predicted by the two cues differs; this mechanism could not easily explain the ability of a conditioned inhibitor to reduce behaviors directed at obtaining a specific food involving a new action (but see Russek

et al., 2017). It has recently been shown that the lateral habenula, which is thought to provide input critical to the generation of negative prediction errors by VTA dopamine neurons (Ji and Shepard, 2007; Jhou et al., 2009; Hong et al., 2011), is necessary for outcome-specific conditioned reinforcement (Laurent et al., 2017). Given evidence that transient activation of the dopamine neurons seems to support learning about specific features of impending rewarding (Keiflin et al., 2017) and even nonrewarding events (Sharpe et al., 2017), it will be of interest to test whether the same is true of transient inhibition of these neurons.

References

- Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47:129–141. [CrossRef Medline](#)
- Bayer HM, Lau B, Glimcher PW (2007) Statistics of midbrain dopamine neuron spike trains in the awake primate. *J Neurophysiol* 98:1428–1439. [CrossRef Medline](#)
- Bromberg-Martin ES, Hikosaka O (2009) Midbrain dopamine neurons signal preference for reward information about upcoming rewards. *Neuron* 63:119–126. [CrossRef Medline](#)
- Bromberg-Martin ES, Matsumoto M, Hikosaka O (2010) Dopamine in motivational control: rewarding, aversive and alerting. *Neuron* 68:815–834. [CrossRef Medline](#)
- Chang CY, Esber GR, Marrero-Garcia Y, Yau HJ, Bonci A, Schoenbaum G (2016) Brief optogenetic inhibition of VTA dopamine neurons mimics the effects of endogenous negative prediction errors during pavlovian over-expectation. *Nat Neurosci* 19:111–116. [CrossRef Medline](#)
- Chang CY, Gardner M, Di Tillio MG, Schoenbaum G (2017) Optogenetic blockade of dopamine transients prevents learning induced by changes in reward features. *Curr Biol* 27:3480–3486.e3. [CrossRef Medline](#)
- Chuong AS, Miri ML, Busskamp V, Matthews GA, Acker LC, Sørensen AT, Young A, Klapoetke NC, Henninger MA, Kodandaramaiah SB, Ogawa M, Ramanlal SB, Bandler RC, Allen BD, Forest CR, Chow BY, Han X, Lin Y, Tye KM, Roska B, et al (2014) Noninvasive optical inhibition with a red-shifted microbial rhodopsin. *Nat Neurosci* 17:1123–1129. [CrossRef Medline](#)
- D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319:1264–1267. [CrossRef Medline](#)
- Daw ND, Kakade S, Dayan P (2002) Opponent interactions between serotonin and dopamine. *Neural Netw* 15:603–616. [CrossRef Medline](#)
- Day JJ, Roitman MF, Wightman RM, Carelli RM (2007) Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. *Nat Neurosci* 10:1020–1028. [CrossRef Medline](#)
- Dreyer JK, Herrik KF, Berg RW, Hounsgaard JD (2010) Influence of phasic and tonic dopamine release on receptor activation. *J Neurosci* 30:14273–14283. [CrossRef Medline](#)
- Frank MJ, Moustafa AA, Haughey HM, Curran T, Hutchison KE (2007) Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proc Natl Acad Sci U S A* 104:16311–16316. [CrossRef Medline](#)
- Gershman SJ, Schoenbaum G (2017) Rethinking dopamine prediction errors. *BioRxiv*. Available at <https://www.biorxiv.org/content/early/2018/05/21/239731>.
- Glimcher PW (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proc Natl Acad Sci U S A* 108:15647–15654. [CrossRef Medline](#)
- Hamid AA, Pettibone JR, Mabrouk OS, Hetrick VL, Schmidt R, Vander Wee CM, Kennedy RT, Aragona BJ, Berke JD (2016) Mesolimbic dopamine signals the value of work. *Nat Neurosci* 19:117–126. [CrossRef Medline](#)
- Hart AS, Rutledge RB, Glimcher PW, Phillips PE (2014) Phasic dopamine release in the rat nucleus accumbens symmetrically encodes a reward prediction error term. *J Neurosci* 34:698–704. [CrossRef Medline](#)
- Hong S, Jhou TC, Smith M, Saleem KS, Hikosaka O (2011) Negative reward signals from the lateral habenula to dopamine neurons are mediated by rostromedial tegmental nucleus in primates. *J Neurosci* 31:11457–11471. [CrossRef Medline](#)
- Horvitz JC (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* 96:651–656. [CrossRef Medline](#)
- Horvitz JC, Stewart T, Jacobs BL (1997) Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Res* 759:251–258. [CrossRef Medline](#)
- Jhou TC, Fields HL, Baxter MG, Saper CB, Holland PC (2009) The rostromedial tegmental nucleus (RMTg), a GABAergic afferent to midbrain dopamine neurons, encodes aversive stimuli and inhibits motor responses. *Neuron* 61:786–800. [CrossRef Medline](#)
- Ji H, Shepard PD (2007) Lateral habenula stimulation inhibits rat midbrain dopamine neurons through a GABA_A receptor-mediated mechanism. *J Neurosci* 27:6923–6930. [CrossRef Medline](#)
- Kakade S, Dayan P (2002) Dopamine: generalization and bonuses. *Neural Netw* 15:549–559. [CrossRef Medline](#)
- Keiflin R, Pribut HJ, Shah NB, Janak PH (2017) Phasic activation of ventral tegmental, but not substantia nigra, dopamine neurons promotes model-based Pavlovian reward learning. *BioRxiv* 232678.
- Kim KM, Baratta MV, Yang A, Lee D, Boyden ES, Fiorillo CD (2012) Optogenetic mimicry of the transient activation of dopamine neurons by natural reward is sufficient for operant reinforcement. *PLoS One* 7:e33612. [CrossRef Medline](#)
- Langdon AJ, Sharpe MJ, Schoenbaum G, Niv Y (2018) Model-based predictions for dopamine. *Curr Opin Neurobiol* 49:1–7. [CrossRef Medline](#)
- Laurent V, Wong FL, Balleine BW (2017) The lateral habenula and its input to the rostromedial tegmental nucleus mediates outcome-specific conditioned inhibition. *J Neurosci* 37:10932–10942. [CrossRef Medline](#)
- Mackintosh NJ (1975) A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol Rev* 82:276–298. [CrossRef](#)
- Mahn M, Prigge M, Ron S, Levy R, Yizhar O (2016) Biophysical constraints of optogenetic inhibition at presynaptic terminals. *Nat Neurosci* 19:554–556. [CrossRef Medline](#)
- Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* 459:837–841. [CrossRef Medline](#)
- Mirenowicz J, Schultz W (1994) Importance of unpredictability for reward responses in primate dopamine neurons. *J Neurophysiol* 72:1024–1027. [CrossRef Medline](#)
- Pan WX, Schmidt R, Wickens JR, Hyland BI (2005) Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. *J Neurosci* 25:6235–6242. [CrossRef Medline](#)
- Pearce JM, Hall G (1980) A model for pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev* 87:532–552. [CrossRef Medline](#)
- Raimondo JV, Kay L, Ellender TJ, Akerman CJ (2012) Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission. *Nat Neurosci* 15:1102–1104. [CrossRef Medline](#)
- Rescorla RA (1969) Pavlovian conditioned inhibition. *Psychol Bull* 72:77–94. [CrossRef](#)
- Rescorla RA, Holland PC (1977) Associations in pavlovian conditioned inhibition. *Learn Motiv* 8:429–447. [CrossRef](#)
- Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: *Classical conditioning II: current research and theory* (Black AH, Prokasy WF, eds), pp 64–99. New York, NY: Appleton-Century-Crofts.
- Roesch MR, Calu DJ, Schoenbaum G (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nat Neurosci* 10:1615–1624. [CrossRef Medline](#)
- Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND (2017) Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput Biol* 13:e1005768. [CrossRef Medline](#)
- Schultz W (2016) Dopamine reward prediction-error signalling: a two-component response. *Nat Rev Neurosci* 17:183–195. [CrossRef Medline](#)
- Sharpe MJ, Chang CY, Liu MA, Batchelor HM, Mueller LE, Jones JL, Niv Y, Schoenbaum G (2017) Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nat Neurosci* 20:735–742. [CrossRef Medline](#)
- Steinberg EE, Keiflin R, Boivin JR, Witten IB, Deisseroth K, Janak PH (2013) A causal link between prediction errors, dopamine neurons and learning. *Nat Neurosci* 16:966–973. [CrossRef Medline](#)

- Stopper CM, Tse MTL, Montes DR, Wiedman CR, Floresco SB (2014) Overriding phasic dopamine signals redirects action selection during risk/reward decision making. *Neuron* 84:177–189. [CrossRef Medline](#)
- Sutton RS (1988) Learning to predict by the method of temporal difference. *Machine Learn* 3:9–44. [CrossRef](#)
- Takahashi YK, Batchelor HM, Liu B, Khanna A, Morales M, Schoenbaum G (2017) Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. *Neuron* 95:1395–1405.e3. [CrossRef Medline](#)
- Tobler PN, Dickinson A, Schultz W (2003) Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm. *J Neurosci* 23:10402–10410. [CrossRef Medline](#)
- Tsai HC, Zhang F, Adamantidis A, Stuber GD, Bonci A, de Lecea L, Deisseroth K (2009) Phasic firing in dopamine neurons is sufficient for behavioral conditioning. *Science* 324:1080–1084. [CrossRef Medline](#)
- Waelti P, Dickinson A, Schultz W (2001) Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412:43–48. [CrossRef Medline](#)
- Wiegert JS, Mahn M, Prigge M, Printz Y, Yizhar O (2017) Silencing neurons: tools, applications, and experimental constraints. *Neuron* 95:504–529. [CrossRef Medline](#)
- Witten IB, Steinberg EE, Lee SY, Davidson TJ, Zalocusky KA, Brodsky M, Yizhar O, Cho SL, Gong S, Ramakrishnan C, Stuber GD, Tye KM, Janak PH, Deisseroth K (2011) Recombinase-driver rat lines: tools, techniques, and optogenetic application to dopamine-mediated reinforcement. *Neuron* 72:721–733. [CrossRef Medline](#)
- Zweifel LS, Parker JG, Lobb CJ, Rainwater A, Wall VZ, Fadok JP, Darvas M, Kim MJ, Mizumori SJ, Paladini CA, Phillips PE, Palmiter RD (2009) Disruption of NMDAR-dependent burst firing by dopamine neurons provides selective assessment of phasic dopamine-dependent behavior. *Proc Natl Acad Sci U S A* 106:7281–7288. [CrossRef Medline](#)